

In-Class Exercise: Alternative Credit Scoring Business Models

Exercise 1: Structured Debate — “Is Upstart a Scoring Vendor or a Lender?”

Format: Split into two teams. Each team prepares three arguments for its assigned position, one concession acknowledging the strongest counter-argument, and a closing that addresses the concession. Then the class votes after both sides present.

Team A — “Upstart Is a Scoring Vendor”

Anchoring evidence: Upstart’s core product is a machine-learning underwriting model licensed to partner banks and credit unions. Its per-decision licence, model-risk deliverables, and fair-lending test packs are consumed by lenders who then make their own credit decisions. The model is the product.

Team A: Upstart Is a Scoring Vendor

Argument I

Argument II

Argument III

 Concession *Strongest argument AGAINST your position:*

 Closing *How you address the concession:*

Team B — “Upstart Is a Lender”

Anchoring evidence: Upstart originates loans on its own platform, retains economic exposure to the flow through capital-markets arrangements, and earns a material share of revenue from transactional economics on loans rather than pure model-licensing fees. Its business model cannot be disentangled from the flow of credit itself.

Team B: Upstart Is a Lender

Argument I

Argument II

Argument III

 Concession *Strongest argument AGAINST your position:*

 Closing *How you address the concession:*


Debrief Questions

Q1: Does the answer — scoring vendor or lender — matter for how regulators should supervise Upstart? Why or why not?

- Q2:** Could the answer genuinely be “both”? If so, what does that imply about how to classify hybrid model-plus-flow businesses?
- Q3:** Name another business (in any sector) that blurs the line between “provides a decision engine” and “carries the economic risk of the decisions”. What tensions does that blurring create?

Exercise 2: Credit Value Chain Mapping

Scenario: The credit value chain can be broken into six links. Alternative-credit-scoring vendors attack specific links, sometimes creating new links that did not previously exist. Your task: for each link, identify a vendor from the reference slate, describe the value the vendor creates, and classify the outcome.



Value Chain Link	Vendor tacking It	At- Value Created	New Link or Old?	Moat or Margin Cap?
Consumer Consent				
Data Enrichment				
Score Construction				
Decisioning				
Reason Codes				
Monitoring				

Synthesis Questions

- Q1:** Which link in the credit value chain creates the strongest moat for a scoring vendor, and which link most clearly caps its margin? Defend both answers with reference to label accumulation, regulatory barriers, and data ownership.
- Q2:** For at least two of the five vendors on the reference slate, identify a link where the vendor rents rather than owns the underlying function. What would it take for the vendor to convert that rented link into an owned one, and should it?

Facilitator Solutions

Sample answers for instructor reference. These are illustrative; student reasoning may diverge and still be valid.

Exercise 1: Debate Sample Answers

Team A (Upstart Is a Scoring Vendor) — sample arguments

Argument I. Upstart's core deliverable to each partner lender is a machine-learning underwriting decision, not a loan. The Business Model Canvas reads as a data-as-input B2B SaaS vendor: the Key Activities are model training, reason-code generation, and fair-lending testing; the Key Resources are the feature pipeline and the labelled-outcome dataset; the Revenue Streams are per-decision licences and model-risk deliverables bundled with each decision. Those blocks describe a vendor, not a lender.

Argument II. Upstart sells into a buyer who makes the credit decision — the partner lender. The lender's compliance team, not Upstart, signs the adverse-action letter and owns the credit policy. That division of responsibility is the hallmark of a vendor relationship: Upstart supplies the decision engine; the lender owns the consequence. A lender-classified firm would own the consequence directly.

Argument III. Upstart's competitive advantage resides in its model IP, its label history, and its fair-lending-testing and reason-code-generation tooling — all assets typical of a B2B SaaS platform. None of those advantages are balance-sheet advantages. A genuine lender's competitive advantage would include funding costs, regulatory capital treatment, and deposit franchise — none of which Upstart competes on primarily.

Concession. The strongest argument against Team A is that Upstart does originate loans on its own platform and retains economic exposure to the flow through capital-markets arrangements, which complicates a pure vendor classification.

Closing. Self-origination was a chicken-and-egg wedge that generated the first labels before partner lenders trusted the score; retained economic exposure on a portion of the flow aligns the vendor's incentives with the partner lender's. Both are structural features of a platform that sells decisions, not loans. The vendor classification therefore describes the economic substance.

Team B (Upstart Is a Lender) — sample arguments

Argument I. Upstart carries real economic exposure on loans that flow through its platform, including self-originated loans held on its own balance sheet and residual exposure on securitisation structures. Those exposures are balance-sheet positions, not software licences. A firm whose profit and loss moves with the credit cycle is a lender by economic substance, regardless of how its Canvas boxes are labelled.

Argument II. The revenue mix at Upstart is dominated by transactional economics on loans that flow through its origination partners and capital-markets activity, not by pure model-licensing fees. If the model-licensing fee alone were stripped out, the remaining business would look like a lender that uses technology to underwrite rather than a technology vendor that happens to have lender customers.

Argument III. Regulators have consistently scrutinised Upstart under lending-focused frameworks — consumer-credit disclosures, fair-lending testing, adverse-action rules — rather than under software-vendor frameworks. The supervisory treatment is calibrated to the credit risk the firm's decisions create, not to the interface through which those decisions are delivered. If the regulator sees a lender, the economic substance is a lender.

Concession. The strongest argument against Team B is that Upstart does not take deposits, does not run a branch network, and earns a material share of revenue per decision through a B2B API — all features of a scoring vendor rather than a traditional lender.

Closing. The absence of deposits does not change the fact that Upstart earns on credit flow and bears credit risk. A lender without a deposit franchise is still a lender — it simply funds itself through capital-markets structures rather than retail deposits. Economic substance over Canvas labels: Upstart is a lender that happens to deliver underwriting through an API.

Debrief Q1 — Regulatory supervision

Whether regulators should supervise Upstart as a lender or a scoring vendor depends on the risks the firm creates rather than on the preferred self-description. If the firm carries credit exposure on its own balance sheet or through retained tranches in securitisations, prudential and consumer-credit rules apply to that exposure directly. If the firm supplies decisions to partner lenders who themselves carry the exposure, the relevant oversight is model-risk governance on the vendor side and lender-side accountability for the decisions used. In practice Upstart occupies both positions simultaneously, which is why supervisors have increasingly applied fair-lending, consumer-credit, and model-risk frameworks in parallel. The classification matters because each framework carries different capital, disclosure, and examination obligations.

Debrief Q2 — “Both” as an answer

The answer genuinely can be “both”. Upstart is a scoring vendor with a portion of the flow carried on its own platform and a portion distributed to lender partners. That duality reveals the limits of categorical labels inherited from a world where underwriting, funding, and origination were bundled inside one institution. When a firm unbundles decision, funding, and origination across a platform architecture, the right regulatory response is functional: supervise each function where it resides, regardless of the firm’s overall label. “Both” as an answer implies that classification by function is more robust than classification by firm.

Debrief Q3 — Cross-sector blurring example

A close parallel is a ride-hailing platform that began as pure software matching drivers to riders and then began carrying its own fleet in some markets. The platform is a decision engine (matching, pricing, routing) plus, in those markets, an operator with economic exposure to the vehicles it deploys. Regulators have struggled to classify it: transportation carrier or software intermediary. The tension is directly analogous to Upstart’s: the firm delivers decisions as software but carries some of the operational risk those decisions create. In both cases, functional regulation — regulating the matching, pricing, or underwriting function separately from the balance-sheet or fleet-ownership function — is better-calibrated than categorical classification.

Exercise 2: Credit Value Chain Mapping Sample Answers

Value Link	Chain	Vendor	Value Created	New or Old Link?	Moat or Margin Cap?
Consumer Consent	Con-	Experian Boost (consumer-permissioned utility and telecom tradeline access)	Creates a consumer-facing consent flow that did not previously exist inside the bureau operating model	New	Moat (consumer relationship is a new asset owned by the bureau)
Data Enrichment	Enrich-	Zest AI (feature engineering across alternative-data sources)	Converts raw alternative-data signals into model-ready features with governance metadata attached	Old	Margin Cap (data providers retain pricing power upstream)
Score Construction	Con-	Upstart (machine-learning underwriting model)	Produces a decision with wider approved envelope than bureau-only scoring	Old	Moat (label accumulation across partner lenders compounds over time)
Decisioning		Pagaya (envelope widening with securitisation exit)	Couples incremental approvals to a capital-markets exit pipe that absorbs the flow	Old	Moat (ABS infrastructure is expensive to replicate; margin cap only where capital-markets depth is thin)
Reason Codes		FairPlay (fair-lending-testing and remediation)	Converts a disparate-outcome test result into an actionable model revision package the lender can sign off	New (the packaged-service form did not previously exist)	Moat if supervisory cadence keeps rising; margin cap if a standard toolkit arrives
Monitoring		Zest AI (drift detection and retraining-as-a-service)	Keeps the model in calibration between re-training cycles and triggers action when drift breaches tolerance	Old	Moat (requires the same labels and governance as the core model, lending a natural bundling advantage)

Synthesis Question 1 Sample Answer

The link that creates the strongest moat for a scoring vendor is Score Construction. Once a vendor has accumulated a label history across many partner lenders, that label asset compounds in a way a new entrant cannot replicate without years of flow. Each additional label sharpens the model and widens the envelope, which in turn attracts more origination partners and deepens the label pool. Regulatory barriers reinforce the moat: model-risk documentation, reason-code generation, and fair-lending test packages become part of the vendor's asset base, all reusable across the label history. The link most likely to cap margin is Data Enrichment. Upstream data providers — bureaus, telecom operators, utility aggregators — retain pricing power, and the vendor is a buyer. As the vendor's flow scales, its data bill scales with it, which caps gross margin unless the vendor can originate its own consented data (as Experian Boost does at the consent link).

Synthesis Question 2 Sample Answer

Upstart rents much of its data enrichment: it consumes bureau feeds and third-party alternative-data aggregators as inputs. Converting those rented links into owned ones would require Upstart to build its own consumer-consent layer — effectively replicating what Experian Boost has done — so that alternative data flows into the model under a consent relationship Upstart directly owns. The case for converting is gross-margin expansion and strategic independence from bureau pricing; the case against is the operational burden of running a consumer-consent product, and the marketing cost of competing for the consumer-facing relationship that Experian Boost already occupies. FairPlay rents its distribution: it reaches lenders through integration partners and consulting channels rather than direct sales. Converting that rented link would mean building an in-house enterprise sales motion targeted at lender compliance and model-risk teams, which is slower but yields higher retention and higher gross margin per account. Whether either vendor should convert is a strategic choice between speed-to-scale (renting) and margin-durability (owning); the right answer depends on how long the vendor expects the wedge to remain open before a standard toolkit or an incumbent bundle appears.